A Trajectory Planning Algorithm for Quantifying Space Ownership in Professional Football

Kyle Burris

January 2019

1 Introduction

American football, like many other popular team sports, is fundamentally about the creation and use of space. Offensive linemen and defensive linemen struggle to control the space near the quarterback. Wide receivers run precise routes to gain separation from defensive backs. Linebackers are given gap assignments to control the running game. Despite the significance of space in the game of football, few methods exist for quantifying player occupation and the creation of space on the football field.

Significant attention has been devoted to understanding the value of space in sports such as basketball and soccer. [1] uses a weighted Voronoi concept, in which a given location on a basketball court is assumed to be owned by the nearest player. The degree of ownership is defined to be inversely proportional to the distance from the nearest player, restricted to the unit interval. While this definition of space ownership is reasonable and easy to implement, it fails to take into account the direction and magnitude of player motion. If a player is running at a full sprint, an ideal space ownership model would assign greater ownership to areas in the direction of motion and less ownership to areas away from the direction of motion.

In an attempt to address this problem and the non-exclusive nature of space ownership in the game of soccer, [2] uses a bivariate normal density function to quantify space creation. The mean vector and covariance matrix of this normal density is a function of the player's initial velocity vector and the player's distance from the ball. This model works well when players are assumed to follow a consistent trajectory over time. However, it is not physically motivated, failing to account for potential sharp changes of direction that characterize the game of football.

In this article, we propose a physical model of space ownership that not only takes initial speed and angle into account, but allows for player-specific variation in deceleration, change-of-direction ability, acceleration, and maximum speed.

The fundamental idea behind our approach is that *a space is owned by the player who can beat every other player to that space*. This has broad applications to the evaluation of quarterbacks and receivers, since this model can identify open receivers well in advance of other models that have been previously proposed.

2 Optimal Player Trajectory Planning

In order to determine which player will arrive at a location first, we need to first determine the time-optimal trajectory for each player and then compute the time that it will take each player to arrive at that location. This type of problem has been heavily studied in the robotics and autonomous vehicles literature [3, 4] and typically involves formulating the problem as a constrained non-linear programming problem. Smoothing splines or B-splines are often used to model the trajectory of agents as well [5].

As is common in the literature, we also treat optimal player trajectory planning as a non-linear programming problem. However, we impose simplifying assumptions on the trajectory to simultaneously improve model interpretability and optimization stability. Specifically, we assume that a player's time-optimal trajectory to a location consists of a maximum of three ordered phases: exponential deceleration, constant-velocity turn, and exponential acceleration. This eliminates all routes in which a player performs multiple changes of direction, which are sub-optimal without the presence of dynamic obstacles. Further, we assume that a player will turn no more than 360 degrees during his optimal route, preventing local optima where a player turns in circles for a period of time. Finally, we assume that there exists a maximum centripetal force that a player's body can endure without injury, for it becomes increasingly difficult for players to make sharp turns as velocity increases. We believe that this bound on centripetal force increases throughout youth player development, as [6] found that young rugby players take more rounded, inefficient turns than older rugby players.

In our model, we assume that player-specific differences in movement time can be attributed to five distinct parameters: maximum controlled deceleration rate, maximum controlled acceleration rate, maximum speed, maximum centripetal force, and player mass. Notationally, we denote these parameters by k_d, k_a, v_{max}, F_{max}, and m respectively. Suppose that a player is moving at speed s_0 in direction θ_0 and is currently at location $p_0 = (x_0, y_0)$. Given that the player's path follows a straight path, followed by a sharp circular cut, followed by another straight path to the destination, we would like to find the minimum amount of time that it will take for the player to reach location $p_1 = (x_1, y_1)$. Our model assumes that the speed and direction of the player's movement from p_0 to p_1 is dictated by the following set of parametric equations:

$$s(t) = \begin{cases} s_0 e^{-k_d t} & t \leq t_1 \\ s(t_1) & t_1 < t \leq t_2 \\ s(t_1) + (v_{max} - s(t_1))(1 - e^{-k_a(t-t_2)}) & t_2 < t \leq t_3 \end{cases} \tag{1}$$

$$\theta(t) = \begin{cases} \theta_0 & 0 \leq t \leq t_1 \\ \theta_0 + \omega(t - t_1) & t_1 < t \leq t_2 \\ \theta_0 + \omega(t_2 - t_1) & t_2 < t \leq t_3, \end{cases} \tag{2}$$

where ω is the player's angular velocity. The position of the player at $t = 0$ is (x_0, y_0) and the position at $t = t_3$ is (x_1, y_1). The change in x from $t = 0$ to $t = t_3$ is $\int_0^{t_3} s(t)\cos(\theta(t))dt$. The change in y from $t = 0$ to $t = t_3$ is $\int_0^{t_3} s(t)\sin(\theta(t))dt$. The free variables in this case are t_1, t_2, t_3, and ω. As such, we have two equality constraints and five inequality constraints:

$$x_1 - x_0 = \int_0^{t_3} s(t)\cos(\theta(t))dt \tag{3}$$

$$y_1 - y_0 = \int_0^{t_3} s(t)\sin(\theta(t))dt \tag{4}$$

$$-2\pi \leq \omega(t_2 - t_1) \leq 2\pi \tag{5}$$

$$t_1 > 0 \qquad t_2 > t_1 \qquad t_3 > t_2 \tag{6}$$

$$-F_{max}/m, \leq s(t_1)\omega \leq F_{max}/m, \tag{7}$$

where m is the mass of the player. Evaluating the integrals in (3) and (4), we have that

$$\begin{aligned} \int_0^{t_3} s(t)\cos(\theta(t))dt &= \int_0^{t_1} s(t)\cos(\theta(t))dt + \int_{t_1}^{t_2} s(t)\cos(\theta(t))dt + \int_{t_2}^{t_3} s(t)\cos(\theta(t))dt \\ &= \int_0^{t_1} s_0 e^{-k_d t}\cos(\theta_0)dt + \int_{t_1}^{t_2} s(t_1)\cos(\theta_0 + \omega(t - t_1))dt \\ &+ \int_{t_2}^{t_3} s(t_1) + (v_{max} - s(t_1))(1 - e^{-k_a(t-t_2)})\cos(\theta(t_2))dt \\ &= \frac{s_0\cos(\theta_0)}{k_d}(1 - e^{-k_d t_1}) + \frac{s(t_1)}{\omega}(\sin(\theta_0 + \omega(t_2 - t_1)) - \sin(\theta_0)) \\ &+ \cos(\theta(t_2))\left[\frac{(s(t_1) - v_{max})}{k_a}(1 - e^{-k_a(t_3-t_2)}) + v_{max}(t_3 - t_2)\right] \end{aligned} \tag{8}$$

$$
\begin{aligned}
\int_0^{t_3} s(t)\sin(\theta(t))dt &= \int_0^{t_1} s(t)\sin(\theta(t))dt + \int_{t_1}^{t_2} s(t)\sin(\theta(t))dt + \int_{t_2}^{t_3} s(t)\sin(\theta(t))dt \\
&= \int_0^{t_1} s_0 e^{-k_d t}\sin(\theta_0)dt + \int_{t_1}^{t_2} s(t_1)\sin(\theta_0 + \omega(t - t_1))dt \\
&+ \int_{t_2}^{t_3} s(t_1) + (v_{max} - s(t_1))(1 - e^{-k_a(t-t_2)})\sin(\theta(t_2))dt \\
&= \frac{s_0\sin(\theta_0)}{k_d}(1 - e^{-k_d t_1}) - \frac{s(t_1)}{\omega}(\cos(\theta_0 + \omega(t_2 - t_1)) - \cos(\theta_0)) \\
&+ \sin(\theta(t_2))\left[\frac{(s(t_1) - v_{max})}{k_a}(1 - e^{-k_a(t_3-t_2)}) + v_{max}(t_3 - t_2)\right]
\end{aligned} \tag{9}
$$

To facilitate ease of computation, we reparameterize the model in terms of $s(t_1)$, $s(t_1)/\omega$, $\omega(t_2 - t_1)$, and $t_3 - t_2$. We then use sequential least squares quadratic programming techniques [7] to find a locally optimal trajectory that satisfies the imposed physical constraints (3) - (7), as implemented in the `nloptr` R package [8]. We find suitable initial values using the L-BFGS-B optimization algorithm, assuming that the turn speed is a tenth of the initial speed and that the turn radius is double the lower bound implied by the maximum centripetal force. The algorithms iterate until convergence. For a small number of initial conditions (less than 1%), the algorithm failed to convergence, so we used random forests trained on all other trajectories to impute these times.

3 Empirical Example

Consider a situation in which a player weighs 200 lbs and whose maximum deceleration rate and maximum acceleration rate are both equal to 2. Further suppose that his maximum speed is 9 meters per second, and that the maximum amount of centripetal force he can endure is equal to 400π pound meters per second squared. These numbers, while intended to be somewhat realistic, are merely for the sake of example. Ideally, teams can estimate these quantities using data collected from the scouting combine, player tracking technology, or via wearable technology.

Given a variety of initial velocities and initial angles from the target, we would like to compare the estimated time it takes this player to reach the target under our model. Figure 1 provides comparisons for this player at distances of 2, 5, 15, and 40 yards away from the target.

From Figure 2, it is clear that players most control the space near them when they are moving slowly. At two yards away, a player moving at top speed with

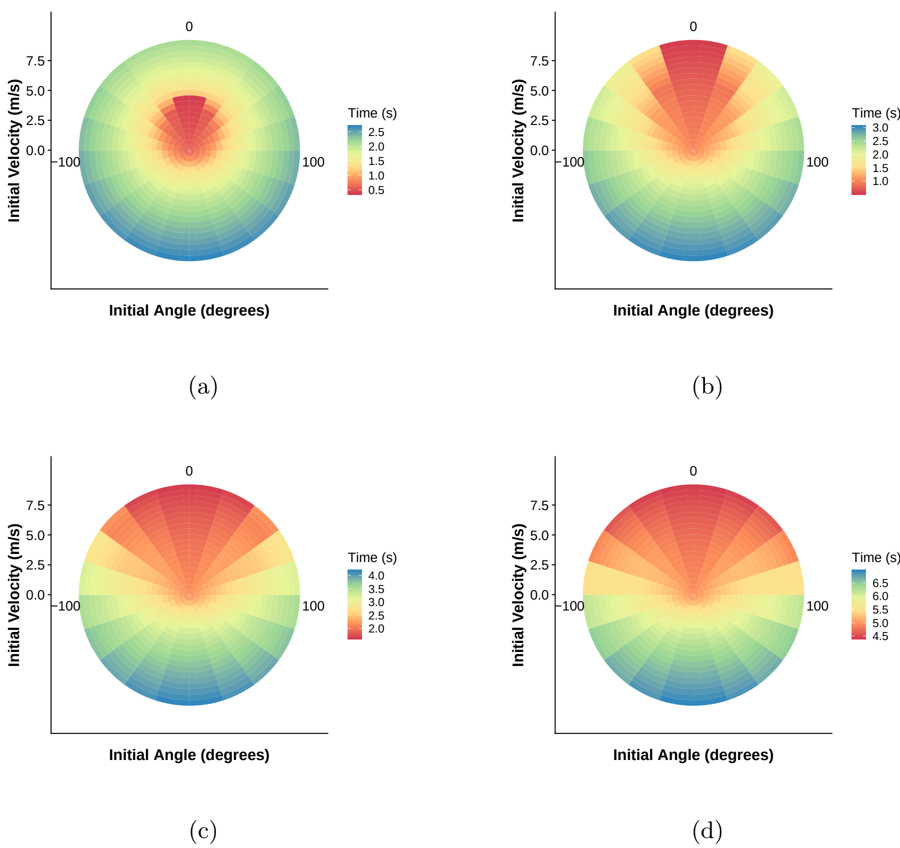

Figure 1: Time to location via optimal trajectory at distance of (a) 2 yards, (b) 5 yards, (c) 15 yards, and (d) 40 yards. A total of 400 velocity/angle combinations were used to make this graph. For the sake of illustration, the player is assumed to be moving with at least a slight angle with the target.

even a small angle from the target will overshoot his intended location, forcing him to backtrack. As the distance from the target increases, it becomes better for a player to be going faster, so long as his initial velocity vector is less than 90 degrees from the target. There is a substantial difference in the time it takes to reach the destination by velocity and angle. The time it takes to decelerate and make a 180 degree turn is over a second.

4 Applications to NFL Player Tracking Data

Player tracking data represents the next-generation of information regarding player performance and movement on the football field. To apply the methodology developed above to this data, we need to use our model to calculate the time that it will take each player to reach all locations on the field, updating that calculation every tenth of a second. This is computationally prohibitive, so we decide to use a neural network to approximate travel times under an optimal trajectory. To do this, we randomly simulate distances from a truncated normal distribution and angles uniformly from $[-\pi, \pi]$. For each of these pairs, we use our physical model to estimate player travel time. We then use this data to train a neural network optimized to minimize the mean-squared error between its predictions and our model outputs. The network has two hidden layers with 64 and 32 nodes respectively, with rectified linear unit activation functions. We also partition the football field into square yards, so as to limit the number of predictions that we will need to make in a given frame.

The mean absolute validation error of our neural network is 0.03 seconds, and the network generally does well in capturing clear instances of space ownership. Using this neural network, we then predict the time that it would take each player to reach every square yard on the field, identifying the player achieving the estimated minimum time for each location. All players are assumed to have the same constants and bounds as the example player in Figure 1.

Consider a play from Week 4 of the 2017 regular season between the Oakland Raiders and the Denver Broncos. Oakland quarterback Derek Carr (Black #4) takes the snap from the shotgun formation and pump fakes to tight end Clive Walford (Black #88), freezing Denver safety Darian Stewart (Orange #26). That's all the time that Oakland receiver Johnny Holton (Black #16) needs to get behind Stewart for a 64-yard catch-and-run touchdown. Figure 2 illustrates the space created during a crucial portion of the play, where Holton beats Stewart over the top.

At the moment of the snap (a), each side largely controls the space on their side of the line of scrimmage. Some receivers control space just past the line of scrimmage due to the lack of press coverage. The Broncos play very good coverage for the first 2.3s of the play (b). However, once Darian Stewart bites on the pump fake (c), Johnny Holton claims a large swath of space in the deep

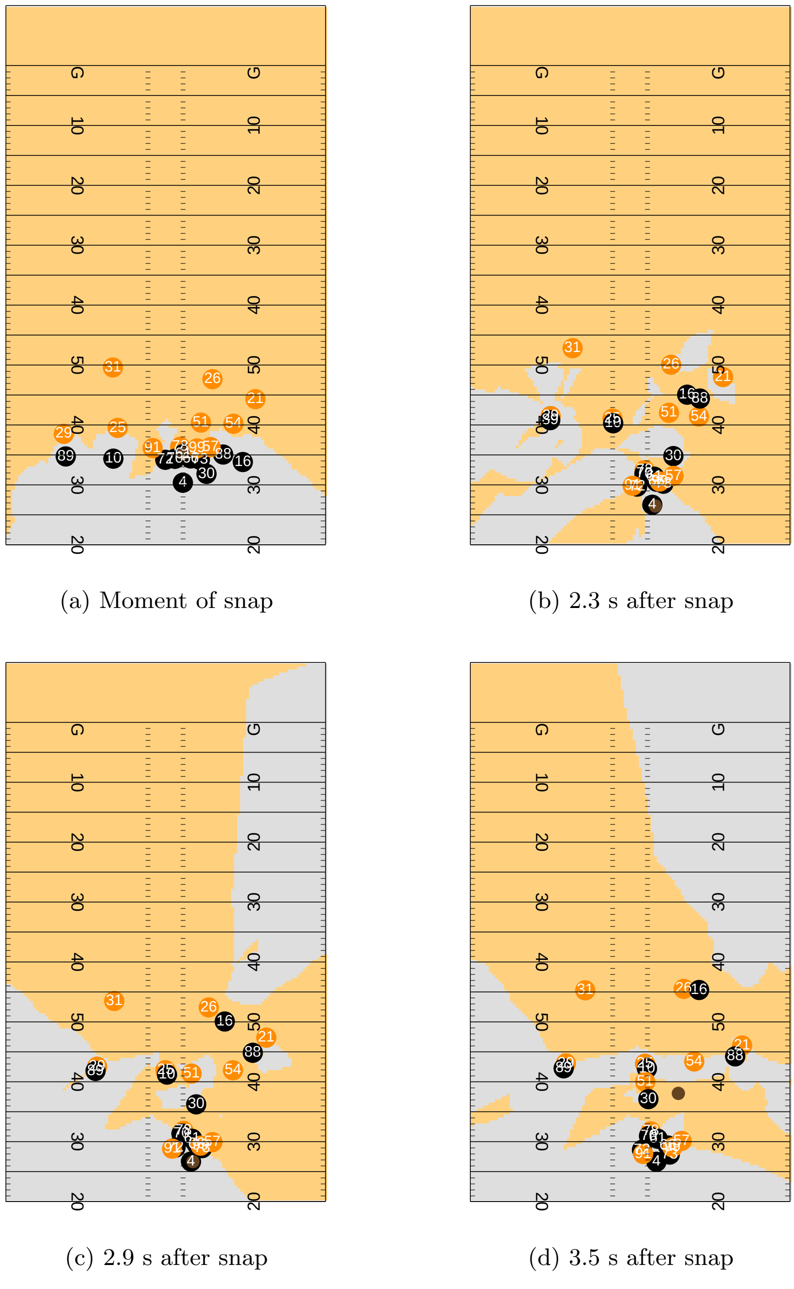

(a) Moment of snap

(b) 2.3 s after snap

(c) 2.9 s after snap

(d) 3.5 s after snap

Figure 2: Neural network approximation to the space ownership model.

right portion of the field. When the ball is released (d), Holton controls nearly the entire right side of the field past the Denver 40 yard line.

5 Discussion

In this paper, we develop a novel, yet general physical model to quantify space ownership in professional sports. This model has particular applications in the game of football, since players change direction quickly in order to create space. This method allows for player-specific variation in speed, acceleration, deceleration, and change-of-direction ability. By incorporating initial conditions such as speed and direction of movement, we are able to identify when a receiver beats a defender well before Voronoi-based models.

There are many potential uses for this space ownership model within professional football. For one, quarterback decision making can be evaluated with more context, since quarterbacks need to anticipate receivers coming open. With this model, we account for receivers slowing down in preparation of a cut, enabling the identification of open receivers much more quickly. Secondly, space ownership models can be leveraged with ball positioning data to create a real-time expected yards metric from a play, similar to the expected possession value introduced in [9] for the game of basketball. The efficacy of coverage and pass rush can also be assessed via a space ownership model, as defensive backs aim to reduce the amount of space owned by the wide receivers and pass rushers attempt to constrict the space that a quarterback has to operate.

Due to the time restrictions associated with this paper, there were many concepts that we were unable to explore in this paper and would like to further develop in the future. In particular, players are not able to freely take any routes to points on the field because they are allowed to be blocked by other players. As such, the current version of the model does not capture space creation near the line of scrimmage very well, since offensive linemen have the ability to control space behind them by blocking pass rushers and defenders are allowed to play press coverage. Future versions may further incorporate adversarial trajectory planning, though this substantially increases the model complexity.

Moreover, we would also like to incorporate a notion of shared ownership of space. As an example, it is difficult to tell who actually owns the space in the left-hand side of Figure 2 (b) since the defender and wide receiver are immediately adjacent to one another. Ownership of a location may be defined to be shared if more than one player is able to catch a ball thrown to that location.

Team sports primarily involve players occupying, creating and taking advantage of space. As more data becomes available in team sports, teams and analysts alike will be able to quantify not only the possession of space, but also the value of space ownership to the team. The methodology developed in this paper has applications to all team sports, not just American football. Nevertheless, a time-optimal trajectory model for player space ownership has broad uses in today's NFL, as teams continue to mine NFL Next Gen Stats data for actionable insights.

Replication code is available at `github.com/burrisk/Big-Data-Bowl`.

References

[1] D. Cervone, L. Bornn, and K. Goldsberry, "NBA court realty," MIT Sloan Sports Analytics Conference, 2016.

[2] J. Fernandez and L. Bornn, "Wide open spaces: a statistical technique for measuring space creation in professional soccer," MIT Sloan Sports Analytics Conference, 2018.

[3] A. Gasparetto and V. Zanotto, "Optimal trajectory planning for industrial robots," *Advances in Engineering Software*, vol. 41, pp. 548–556, 2010.

[4] X. Qian, F. Altche, P. Bender, C. Stiller, and A. de La Fortelle, "Optimal trajectory planning for autonomous driving integrating logical constraints: An MIQP perspective," pp. 205–210, 11 2016.

[5] M. Egerstedt and C. F. Martin, "Optimal trajectory planning and smoothing splines," *Automatica*, vol. 37, no. 7, pp. 1057 – 1064, 2001.

[6] G. Condello, C. Minganti, C. Lupo, C. Benvenuti, D. Pacini, and A. Tessitore, "Evaluation of change-of-direction movements in young rugby players," *International journal of sports physiology and performance*, vol. 8, pp. 52–6, 01 2013.

[7] P. Boggs and J. W. Tolle, "Sequential quadratic programming," *Acta Numerica*, vol. 4, pp. 1–51, 01 1995.

[8] S. G. Johnson, "The nlopt nonlinear-optimization package,"

[9] D. Cervone, A. D'Amour, L. Bornn, and K. Golsberry, "Pointwise: Predicting points and valuing decisions in real time with nba optical tracking data," MIT Sloan Sports Analytics Conference, 2014.